



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Journal of Computational Physics 191 (2003) 542–566

JOURNAL OF
COMPUTATIONAL
PHYSICS

www.elsevier.com/locate/jcp

Higher-order time integration schemes for the unsteady Navier–Stokes equations on unstructured meshes

Giridhar Jothiprasad ^{a,*}, Dimitri J. Mavriplis ^{b,1}, David A. Caughey ^{a,2}

^a *Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853, USA*

^b *National Institute of Aerospace, 144 Research Drive, Hampton, VA 23666, USA*

Received 16 July 2002; received in revised form 14 May 2003; accepted 23 June 2003

Abstract

The efficiency gains obtained using higher-order implicit Runge–Kutta (RK) schemes as compared with the second-order accurate backward difference schemes for the unsteady Navier–Stokes equations are investigated. Three different algorithms for solving the nonlinear system of equations arising at each time step are presented. The first algorithm (nonlinear multigrid, NMG) is a pseudo-time-stepping scheme which employs a nonlinear full approximation storage (FAS) agglomeration multigrid method to accelerate convergence. The other two algorithms are based on inexact Newton's methods. The linear system arising at each Newton step is solved using iterative/Krylov techniques and left preconditioning is used to accelerate convergence of the linear solvers. One of the methods (LMG) uses Richardson's iterative scheme for solving the linear system at each Newton step while the other (PGMRES) uses the generalized minimal residual method. Results demonstrating the relative superiority of these Newton's method based schemes are presented. Efficiency gains as high as 10 are obtained by combining the higher-order time integration schemes such as fourth-order Runge–Kutta (RK64) with the more efficient inexact Newton's method based schemes (LMG).

© 2003 Elsevier B.V. All rights reserved.

Keywords: Jacobian-free Newton; Runge–Kutta methods; Navier–Stokes; Multigrid; Unstructured

1. Introduction

The rapid increase in available computational power over the last decade has enabled higher resolution flow simulations and more widespread use of unstructured grid methods for complex geometries. While much of this effort has been focused on steady-state calculations in the aerodynamics community, the need to accurately predict off-design conditions, which may involve substantial amounts of flow separation, points to the need to efficiently simulate unsteady flow fields. Accurate unsteady flow simulations can easily

* Corresponding author. Graduate Student.

E-mail addresses: gj24@cornell.edu (G. Jothiprasad), dimitri@icase.edu (D.J. Mavriplis), dac5@cornell.edu (D.A. Caughey).

¹ NIA Research Fellow.

² Professor.

require several orders of magnitude more computational effort than a corresponding steady-state simulation. For this reason, techniques for improving the efficiency of unsteady flow simulations are required if such calculations are to be feasible in the foreseeable future. The purpose of this work is to investigate possible reductions in computer time due to the choice of an efficient time integration scheme from a series of schemes differing in the order of time accuracy, and by the use of more efficient techniques to solve the nonlinear equations which arise while using implicit time integration schemes. This investigation is carried out in the context of a two-dimensional unstructured mesh laminar Navier–Stokes solver.

Implicit in any comparison of efficiency is a precise error tolerance requirement. For stringent accuracy requirements, high-order temporal discretization schemes are well known to be superior to lower-order (e.g. second-order) schemes, due to their superior asymptotic properties. However, for many large-scale engineering calculations, such as unsteady Reynolds-averaged Navier–Stokes (RANS) simulations, where the time scales of interest are much larger than the spatial scales, larger temporal error tolerances ($O(10^{-2})$ – $O(10^{-3})$) are generally acceptable, and second-order accurate time discretizations are currently widely used, while higher-order methods are generally avoided due to their increased cost per time step. Recently, the use of higher-order accurate implicit Runge–Kutta (RK) schemes has been shown to produce efficiency gains even for relatively coarse error tolerances using a production structured-mesh Navier–Stokes solver [1].

In this paper, we perform a similar investigation within an unstructured mesh setting. Additionally, we investigate the efficiency of various solution techniques for solving the nonlinear problems which arise at each time step for the various time discretizations considered. We consider three solution techniques, namely, a nonlinear multigrid (NMG) method which solves the nonlinear problem directly through pseudo-time-stepping, and two variants of an inexact Newton scheme, where the linear system at each Newton iteration is partially solved using a linear multigrid (LMG) scheme or a multigrid preconditioned generalized minimal residual (PGMRES) approach. Because high-order time discretizations achieve high temporal accuracy with relatively large time steps, thus increasing the condition number of the nonlinear problem, the use of efficient solvers takes on additional significance in such cases.

NMG methods were originally developed for steady-state fluid flow problems and subsequently adapted to unsteady flow problems [2–4]. Newton-based methods have often been avoided in this context due to the additional memory overheads incurred by such methods and the difficulties in providing reliable initializations for nonlinear convergence. However, Newton-based methods have been shown to offer the potential for higher computational efficiency by avoiding frequent nonlinear residual evaluations [5]. Furthermore, the disadvantages of Newton-based methods are less relevant in the context of an unsteady flow solver, where a close initial solution is always available from the previous time step, and where memory considerations are often secondary to CPU time considerations.

In this paper, we illustrate the potential savings achieved using higher-order time discretizations and more efficient nonlinear solvers for unsteady flow simulations on unstructured grids. We investigate the interaction between the time-discretization scheme and the nonlinear solution technique as a function of temporal accuracy and show that the beneficial effects are multiplicative, producing up to an order of magnitude savings in computational effort.

2. Base solver

2.1. Spatial discretization

For the purpose of comparison, an existing two-dimensional unstructured multigrid steady-state Navier–Stokes solver developed in [6] was modified to simulate transient flows by incorporating various

physical time-stepping schemes. The flow equations are discretized using a finite-volume approach. Flow variables are stored at the vertices of the mesh, and control volumes are formed by the median-dual graph of the original mesh, as shown in Fig. 1. A control-volume flux balance is computed by summing fluxes evaluated along the control-volume faces, using the average values of the flow variables on either side of the face in the flux computation. The construction of convective terms corresponds to a central difference scheme which requires additional dissipation terms for stability. These may either be constructed explicitly as a blend of Laplacian and biharmonic operators, or may be obtained by writing the residual of a standard upwind scheme as the sum of a convective and dissipation term

$$\sum_{k=1}^{\text{neighbors}} \frac{1}{2} \{ \mathbf{F}(\mathbf{w}_i) + \mathbf{F}(\mathbf{w}_k) \} \cdot \mathbf{n}_{ik} - \frac{1}{2} |A_{ik}| (\mathbf{w}_L - \mathbf{w}_R), \quad (1)$$

where the convective fluxes are denoted by $\mathbf{F}(\mathbf{w})$, \mathbf{n}_{ik} represents the normal vector of the control-volume face separating the neighboring vertices i and k , and A_{ik} is the flux Jacobian evaluated in the direction normal to this face. The variables \mathbf{w}_L and \mathbf{w}_R represent extrapolated flow properties at the left- and right-hand sides of the control-volume face, respectively. A matrix-based artificial dissipation scheme is obtained by utilizing the same transformation matrix $|A_{ik}|$ as the upwind scheme, but using this to multiply a difference of blended first and second differences rather than a difference of reconstructed states at control-volume boundaries. For the calculations performed in this work, which involve only subsonic flows, the matrix dissipation formed using only second differences has been used exclusively, and the physical viscous terms for the Navier–Stokes equations are discretized to second-order accuracy using a finite-volume approximation.

2.2. Temporal discretization

Time is discretized in a fully implicit sense using both multistep backward difference formulas (BDF) and multistage RK schemes. There are two mathematical properties that are desirable of a numerical integrator [1,7,8]. The first is the “A-stability” property which guarantees that all eigenvalues lying in the left-half of the complex plane will have an amplification factor not greater than unity, independent of the chosen step size. Hence, the only restriction on the time step with an A-stable scheme is the consideration of solution accuracy. The second is the “L-stable” property which guarantees that eigenvalues approaching $-\infty$ are damped in one time step.

Multistep BDF formulas, and in particular the second-order accurate BDF scheme (BDF2), are widely used in the computation of large-scale engineering flows. These schemes require only one nonlinear set of




Fig. 1. Median control volumes for triangular meshes.

equations to be solved at each time step. They suffer, however, from not being self-starting, are difficult to use with variable time steps, and are not A-stable beyond second-order temporal accuracy.

On the other hand, multistage RK schemes are self-starting, are easily implemented in a variable time-stepping mode, and can be designed with A- and L-stability properties for any temporal order. However, these schemes require multiple nonlinear solves at each time step, and hence have often been discounted as noncompetitive compared to BDF schemes. One of the objectives of this paper is to investigate the relative efficiencies of BDF and RK schemes in computing time-dependent solutions to a given level of accuracy.

Consider the integration of the system of ordinary differential equations (ODEs) represented by the equation

$$\frac{d\mathbf{w}}{dt} = \mathbf{S}(t, \mathbf{w}(t)), \tag{2}$$

where the vector \mathbf{S} results from the spatial discretization of the equations of fluid mechanics. The general formula for a k -step BDF scheme can be written as

$$\mathbf{w}^{n+k} = - \sum_{i=0}^{k-1} \alpha_i \mathbf{w}^{n+i} + \Delta t \beta_k \mathbf{S}^{n+k}. \tag{3}$$

BDF schemes require the storage of $k + 1$ solution levels and the computation of one nonlinear solution at each time step. For $k = 2$, the BDF2 scheme is obtained using the coefficients: $\alpha_0 = 1/3$, $\alpha_1 = -4/3$, $\beta_2 = 2/3$. More details of these standard schemes can be found in [1,9].

Runge–Kutta methods are multistage schemes and are implemented as

$$\mathbf{w}^k = \mathbf{w}^n + (\Delta t) \sum_{j=1}^s a_{kj} \mathbf{S}(\mathbf{w}^j), \quad k = 1, 2, \dots, s, \tag{4}$$

$$\mathbf{w}^{n+1} = \mathbf{w}^n + (\Delta t) \sum_{j=1}^s b_j \mathbf{S}(\mathbf{w}^j), \tag{5}$$

where s is the number of stages and a_{ij} and b_j are the Butcher coefficients of the scheme. Following the previous work by Bijl et al. [1] we focus on the ESDIRK class of RK schemes, which stands for explicit first stage, single diagonal coefficient, diagonally implicit Runge–Kutta. The Butcher table for a six-stage ESDIRK scheme is shown in Table 1.

In Table 1, the c_k coefficients, which are of the form: $c_1 = 0$, $c_s = 1.0$, $c_{k-1} < c_k < c_{k+1}$, denote the point in the time interval $[t, t + \Delta t]$. Although these coefficients are not used in the evaluation of the time-stepping scheme, they are shown to illustrate the fact that the RK scheme involves the evaluation of the residual at various locations in time for each time step. These schemes are characterized by a lower triangular form of the coefficient table, thus resulting in a single implicit solve at each individual stage. The first stage is explicit

Table 1
Butcher tableau for the ESDIRK class of RK schemes with number of stages, $s = 6$

$c_1 = 0$	0	0	0	0	0	0
c_2	a_{21}	a_{66}	0	0	0	0
c_3	a_{31}	a_{32}	a_{66}	0	0	0
c_4	a_{41}	a_{42}	a_{43}	a_{66}	0	0
c_5	a_{51}	a_{52}	a_{53}	a_{54}	a_{66}	0
$c_6 = 1$	b_1	b_2	b_3	b_4	b_5	a_{66}
\mathbf{w}^{n+1}	b_1	b_2	b_3	b_4	b_5	a_{66}

($a_{k1} = 0$) and the last stage coefficients take on the form $a_{kj} = b_j$, thus enabling Eq. (5) to be simplified as

$$\mathbf{w}^{n+1} = \mathbf{w}^{k=s}. \quad (6)$$

We use the following notation, RK xy refers to an ESDIRK scheme which has x stages and y th order accuracy. Bijl et al. [1] have compared these schemes and found RK64 to perform well. The numerical values for the coefficients of this scheme are given in Appendix A. More details in general about ESDIRK schemes can be found in [8].

3. Implicit solution technique

Both BDF and RK xy schemes require the solution of a nonlinear system of equations. BDF schemes require the solution of one nonlinear equation per time step. In the case of BDF, a nonlinear residual, $\mathbf{R}(\mathbf{w})$, can be defined from Eq. (3) and is given by

$$\mathbf{R}(\mathbf{w}) \equiv \mathbf{R}(\mathbf{w}^{n+k}) \equiv \frac{\mathbf{w}^{n+k}}{\Delta t} - \beta_k \mathbf{S}^{n+k} + \mathbf{SRC}_{\text{BDF}}, \quad (7)$$

where the superscript on \mathbf{w} has been dropped for the sake of simplicity. $\mathbf{SRC}_{\text{BDF}}$ is the source term independent of $\mathbf{w} \equiv \mathbf{w}^{n+k}$ and is given by

$$\mathbf{SRC}_{\text{BDF}} \equiv \frac{1}{\Delta t} \left[\sum_{i=0}^{k-1} \alpha_i \mathbf{w}^{n+i} \right]. \quad (8)$$

In the case of RK xy schemes, a nonlinear system arises at each stage of the time-stepping scheme and hence more than one nonlinear solve per time step is required. Again, a nonlinear residual, $\mathbf{R}(\mathbf{w})$ for each stage of the RK xy scheme can be defined using Eq. (4) as follows:

$$\mathbf{R}(\mathbf{w}) \equiv \mathbf{R}(\mathbf{w}^k) \equiv \frac{\mathbf{w}^k}{\Delta t} - a_{kk} \mathbf{S}(\mathbf{w}^k) + \mathbf{SRC}_{\text{RK}}, \quad (9)$$

where the superscript on \mathbf{w} has again been dropped for the sake of simplicity. Also, \mathbf{SRC}_{RK} , is the source term independent of $\mathbf{w} \equiv \mathbf{w}^k$ and is given by

$$\mathbf{SRC}_{\text{RK}} \equiv -\frac{\mathbf{w}^n}{\Delta t} - \sum_{j=1}^{k-1} a_{kj} \mathbf{S}(\mathbf{w}^j). \quad (10)$$

Hence, in both BDF and RK xy we are required to obtain the solution of the nonlinear system of equations

$$\mathbf{R}(\mathbf{w}) = 0. \quad (11)$$

Three different methods are proposed for solving Eq. (11) and their relative performances are studied. The three methods, in this paper, are henceforth referred to as:

1. Nonlinear multigrid.
2. Linear multigrid.
3. Preconditioned generalized minimal residual.

In NMG, a pseudo-time-stepping scheme is employed to obtain the solution of the nonlinear system of equations, which is accelerated using a nonlinear full approximation storage (FAS) agglomeration multi-grid method [6,5].

In the other two approaches, an inexact Newton solution strategy is used to solve the nonlinear system of equations [10,6,11]. The resulting linear system of equations is solved using iterative/Krylov techniques. To accelerate convergence, the linear system is left preconditioned using an approximate inverse to the first-order accurate Jacobian which in itself is employed as an approximation to the Jacobian of the second-order accurate discretization. The last two approaches differ only in the methods used to solve the preconditioned linear system of equations. LMG uses the Richardson's iterative method while PGMRES uses the generalized minimal residual method developed by Yaad and Schultz [12].

3.1. Nonlinear multigrid

In a pseudo-time-stepping scheme, the equations are integrated in pseudo-time until Eq. (11) is satisfied.

$$\frac{d\mathbf{w}}{dt^*} + \mathbf{R}(\mathbf{w}) = 0, \quad (12)$$

where t^* is the pseudo-time. Since we do not require a pseudo-time accurate solution, the equations are preconditioned to accelerate the convergence. Hence, we have

$$\mathbf{P}^{-1} \frac{d\mathbf{w}}{dt^*} + \mathbf{R}(\mathbf{w}) = 0. \quad (13)$$

The equations are integrated in pseudo-time using an explicit preconditioned multi-stage scheme [13], which can be written as

$$\begin{aligned} \mathbf{w}^{(0)} &= \mathbf{w}^{[k]}, \\ \mathbf{w}^{(1)} &= \mathbf{w}^{(0)} - \Delta t^* \mathbf{P} \mathbf{R}(\mathbf{w}^{(0)}), \\ &\vdots \\ \mathbf{w}^{(q)} &= \mathbf{w}^{(q-1)} - \Delta t^* \mathbf{P} \mathbf{R}(\mathbf{w}^{(q-1)}), \\ \mathbf{w}^{[k+1]} &= \mathbf{w}^{(q)}, \end{aligned} \quad (14)$$

for a q -stage scheme.

An agglomeration multigrid strategy previously developed as a solver for steady-state problems [13] is used to further accelerate convergence. The agglomeration multigrid may be viewed as a simplified algebraic multigrid strategy. Coarse level grids are constructed by fusing together or agglomerating neighboring control volumes to form a coarser set of larger but more complex control volumes. In the algebraic interpretation of agglomeration multigrid, the coarse levels are no longer geometric grids, but represent groupings of fine grid equations which are summed together to form the coarse grid equations sets [14,15]. The basic smoother on each grid level is a three-stage explicit preconditioned multi-stage scheme with stage coefficients optimized for high frequency damping properties [16].

The preconditioner, \mathbf{P} is chosen to be the block diagonal of the Jacobian of the residual, $\mathbf{R}(\mathbf{w})$,

$$\mathbf{P}_{ii}^{-1} = \mathbf{D}_{ii} = \left[\frac{\partial \mathbf{R}_i}{\partial \mathbf{w}_i} \right]_{\mathbf{w}^{(0)}}. \quad (15)$$

This is referred to as Jacobi preconditioning and provides substantial increases in efficiency when upwind or matrix dissipation discretizations are used [17]. It should be noted that the use of Jacobi preconditioning involves inverting a 4×4 matrix for each vertex at each stage.

3.2. Inexact Newton's methods

To solve the nonlinear system of equations $\mathbf{R}(\mathbf{w}) = 0$, Newton's method requires the solution of a series of linear systems of the form

$$\left[\frac{\partial \mathbf{R}}{\partial \mathbf{w}} \right]_{\mathbf{w}^{[k]}} \delta \mathbf{w}^{[k]} = -\mathbf{R}(\mathbf{w}^{[k]}), \quad (16)$$

where

$$\mathbf{w}^{[k+1]} = \mathbf{w}^{[k]} + \delta \mathbf{w}^{[k]}. \quad (17)$$

Let

$$\mathbf{x} \equiv \delta \mathbf{w}^{[k]}, \quad \mathbf{r} \equiv \mathbf{R}(\mathbf{w}^{[k]}), \quad \mathbf{A} \equiv \left[\frac{\partial \mathbf{R}}{\partial \mathbf{w}} \right]_{\mathbf{w}^{[k]}}. \quad (18)$$

Hence, Eq. (16) now becomes

$$\mathbf{A}\mathbf{x} = -\mathbf{r}. \quad (19)$$

Traditionally, there have been two main obstacles to the use of Newton's method for large-scale multi-physics applications.

(1) *An initial guess inside of the radius of convergence is required for Newton's method to converge.* However, for unsteady problems, a good initial guess is provided by the solution at the previous time step. If the Newton's method does not converge, then by lowering the time step one can get the initial guess as close as necessary to the solution at the next time level. In the calculations presented in this paper, no difficulties were encountered in the convergence of the Newton iterations for any of the time steps used.

(2) *Construction and storage of the Jacobian matrix, A becomes prohibitive.* This problem is particularly exacerbated in three-dimensional for higher-order spatial discretizations which are not confined to the nearest neighbor stencils. This problem is overcome by the use of Jacobian-free methods to solve the linear system of equations. However, additional memory is still required to store the first-order Jacobian for the preconditioning operation, and the various Krylov vectors for the GMRES scheme. On the other hand, the use of additional memory can be rationalized if this produces substantial gains in CPU time, particularly for unsteady flow simulations where cpu time is the dominant concern.

Additionally, in order to improve the computational efficiency of these methods, we use an inexact Newton's method [10,18] where the arising system of equations are not solved exactly. In this paper, we employ a very simplistic method where the number of iterations carried out by the underlying iterative linear solver is held fixed. In its exact form, Newton's method provides quadratic convergence. However due to the various approximations employed by the solution method in this paper, this rate of convergence is not achieved.

3.3. Preconditioned inexact Newton's method

In order to achieve rapid convergence of the linear problem at each Newton iteration, preconditioning methods are used to cluster the eigenvalues of the system. We adopt the approach of left preconditioning in order to achieve this desirable distribution of eigenvalues. The preconditioned system can be written as

$$\mathcal{P}_N \mathbf{A} \mathbf{x} = -\mathcal{P}_N \mathbf{r}. \quad (20)$$

We make the following comments on the preconditioning:

- The preconditioner, $\mathcal{P}_{\mathcal{N}}$, is looked upon here as an operator as opposed to a matrix. Hence, $\mathcal{P}_{\mathcal{N}}$, may or may not be able to be written as a matrix.
- The preconditioner must be chosen as close as possible to A^{-1} so that $\mathcal{P}_{\mathcal{N}}A \approx \mathcal{I}$, where \mathcal{I} is the identity operator.
- Since each step of the Newton’s method, Eq. (16), moves $\mathbf{w}^{[k]}$, towards the solution, \mathbf{w}^* , of $\mathbf{R}(\mathbf{w}) = 0$, any operator which produces a correction $\delta\mathbf{w}^{[k]} = \mathbf{x}$ to advance $\mathbf{w}^{[k]}$ towards \mathbf{w}^* would serve as a reasonable preconditioner.
- Based on the above fact, single or multiple cycles of the NMG method discussed in the earlier subsection could be used as a preconditioner. However, this is computationally inefficient as it does not recognize the fact that we are now solving a system of linear equations.

Keeping in mind the above observations we now propose a better preconditioner, $\mathcal{P}_{\mathcal{N}}$. We first choose $\mathcal{P}_{\mathcal{N}} = \tilde{\mathbf{A}}^{-1}$, where $\tilde{\mathbf{A}}$ is a matrix approximation to the Jacobian, A . Considerations governing the choice of $\tilde{\mathbf{A}}$ would be:

1. *Storage requirements for $\tilde{\mathbf{A}}$ must not be prohibitive.* Storage for $\tilde{\mathbf{A}}$ must use less space than that needed for A or else there would be no space gain in using Jacobian-free methods.
2. *The inverse of $\tilde{\mathbf{A}}$ must be simple to calculate or approximate.* If one is able to compute an approximate inverse to $\tilde{\mathbf{A}}$ fairly easily using iterative methods, this new operator, $\mathcal{P}_{\mathcal{N}} = \tilde{\mathcal{A}}^{-1}$ would serve as an appropriate preconditioner. The two tildes are used to symbolize the fact that there are two approximations involved in the definition of the $\mathcal{P}_{\mathcal{N}}$:
 - (a) $\tilde{\mathbf{A}}$ which is an approximation to the Jacobian, A .
 - (b) An approximation in computing the inverse of $\tilde{\mathbf{A}}$.

In this paper, we choose $\tilde{\mathbf{A}}$ to be the Jacobian of the first-order accurate, nearest neighbor discretization of the nonlinear set of equations. Hence, the storage of $\tilde{\mathbf{A}}$ requires substantially less memory than that of the full Jacobian A of the second-order accurate scheme.

$\tilde{\mathbf{A}}^{-1}$ is computed approximately using a fixed number of W-cycles of a LMG method. This particular multigrid method can be viewed as the linear analogue of the nonlinear FAS agglomeration multigrid scheme described in the NMG method. This approach has been previously described in detail in [5]. In this particular approach the coarse level approximations to the Jacobian are obtained taking the Jacobian of the Galerkin projection of the (frozen) fine grid operator as

$$A_H = \frac{\partial}{\partial w_H} (I_h^H \mathbf{R}_h I_H^h), \tag{21}$$

as opposed to the more traditional LMG Galerkin projection [19]

$$A_H = I_h^H \frac{\partial \mathbf{R}_h}{\partial w_H} I_H^h = I_h^H A_h I_H^h, \tag{22}$$

where I_h^H is the restriction operator and I_H^h is the prolongation operator. This approach was chosen purely for convenience, as the terms in Eq. (21) are readily available. The smoother on each grid was taken as a block diagonal Jacobi solver.

3.4. Linear multigrid

In the method referred to as LMG, the LMG preconditioned system arising at each Newton iteration is solved using a single Richardson iteration. In order to solve the system

$$\mathcal{P}_{\mathcal{N}} \mathbf{A} \mathbf{x} = -\mathcal{P}_{\mathcal{N}} \mathbf{r}, \tag{23}$$

we define the splitting [20]

$$\mathcal{P}_{\mathcal{N}}\mathbf{A} = \mathcal{I} + \mathcal{N}. \quad (24)$$

The resulting iterative scheme is defined as

$$\begin{aligned} \mathcal{I}\mathbf{x}^{(m+1)} &= -\mathcal{P}_{\mathcal{N}}\mathbf{r} - \mathcal{N}\mathbf{x}^{(m)}, \\ \mathcal{I}\delta\mathbf{x}^{(m)} &= -\mathcal{P}_{\mathcal{N}}\mathbf{r} - \mathcal{P}_{\mathcal{N}}\mathbf{A}\mathbf{x}^{(m)}. \end{aligned} \quad (25)$$

As indicated earlier, since we are required to solve Eq. (23) only approximately, we carry out only a single Richardson's iteration. Assuming $\mathbf{x}^{(1)} = \mathbf{0}$, Eq. (25) reduces to

$$\begin{aligned} \delta\mathbf{x}^{(1)} &= -\mathcal{P}_{\mathcal{N}}\mathbf{r}, \\ \mathbf{x}^{(2)} &= \mathbf{x}^{(1)} + \delta\mathbf{x}^{(1)} = \delta\mathbf{x}^{(1)} = -\mathcal{P}_{\mathcal{N}}\mathbf{r}. \end{aligned} \quad (26)$$

Hence, we have

$$\delta\mathbf{w}^{[k]} = \mathbf{x}^{(2)} = -\mathcal{P}_{\mathcal{N}}\mathbf{r}. \quad (27)$$

Eq. (27) illustrates the correspondence of this scheme to a Newton's method in which the first-order accurate Jacobian is used along with a second-order accurate residual.

3.5. Preconditioned generalized minimal residual

Having presented the LMG scheme as a preconditioned Richardson iteration, the PGMRES scheme can similarly be described as the equivalent scheme obtained when the single Richardson iteration is replaced by a GMRES Krylov subspace iterative approach [21]. In this method, we use GMRES to solve Eq. (20) in a matrix-free Newton–Krylov fashion [22], making use of the same preconditioner, $\mathcal{P}_{\mathcal{N}}$, which is used in the LMG method. The matrix-free implementation of PGMRES requires the computation of the product, $\mathcal{P}_{\mathcal{N}}\mathbf{A}\mathbf{x}$, which is approximated using a first-order Taylor series expansion as

$$\mathcal{P}_{\mathcal{N}}\mathbf{A}\mathbf{x} = \frac{\mathcal{P}_{\mathcal{N}}\mathbf{R}(\mathbf{w}^{[k]} + \epsilon\mathbf{x}) - \mathcal{P}_{\mathcal{N}}\mathbf{R}(\mathbf{w}^{[k]})}{\epsilon} = \frac{\mathcal{P}_{\mathcal{N}}\mathbf{R}(\mathbf{w}^{[k]} + \epsilon\mathbf{x}) - \mathcal{P}_{\mathcal{N}}\mathbf{r}}{\epsilon}, \quad (28)$$

where \mathbf{x} is some unit vector and ϵ is a number chosen proportional to the norm of \mathbf{x} times the square root of machine round-off error, close to the square root of machine round-off [10]. We use the restarted form of GMRES with a fixed number of search directions. While increasing the number of Krylov vectors accelerates convergence, storage and cpu time increase with the number of search directions. The optimal number of search directions is therefore determined experimentally.

4. Validation of the temporal scheme

Numerical experiments have been performed to determine the observed order of accuracy of the various time integration schemes. The test problem chosen for this purpose consists of the unsteady laminar flow around a two-dimensional circular cylinder at a Mach number of 0.2. Two different Reynolds numbers, $Re = 1200$ and $Re = 100$ were considered. The initial flow is symmetric with zero lift. As the wake behind the cylinder starts to grow, it becomes unstable and begins to shed vortices from alternate sides of the cylinder. The computational grid is shown in Fig. 2. The mesh contains a total of 19,012 points, with 196 points equally spaced on the surface of the cylinder, and a normal spacing of 10^3 cylinder diameters. The far-field boundary is a circle concentric with the cylinder and diameter given by $40D$ where D is the diameter of the cylinder. A close up of the mesh around the cylinder is shown in Fig. 3




Fig. 2. Computational mesh for circular cylinder.




Fig. 3. Computational mesh in the region around the cylinder.

Time is nondimensionalized as Ut/D where U is the free stream velocity and D is the diameter of the cylinder. The initial condition for the various studies was obtained by simulating the limit cycle behavior for about 10 shedding cycles using a relatively small time step.

4.1. $Re = 1200$ case

Using $\Delta t = 0.025$, the Strouhal number, St , was calculated to be 0.2469. This value is compared with a \sqrt{Re} -formula fit [23] to the $St-Re$ curve obtained from two different data sources:

- “Numerical” laminar flow, two-dimensional computations by Henderson [24]: the curve fit is given by

$$St = \left(0.2731 - \frac{1.1129}{\sqrt{Re}} + \frac{0.4821}{Re} \right),$$

which yields $St = 0.2409$ at $Re = 1200$. It can be seen that this compares very well with our calculations.

- Experimental studies in the wake of a circular cylinder: beyond $Re = 260$, the curve fit is given by

$$St = \left(0.2234 - \frac{0.3490}{\sqrt{Re}} \right).$$

This gives a $St = 0.2133$ which is different from our computed St of 0.2469. This is because the wake of the cylinder transitions to a three-dimensional shedding mode which involves finer-scale streamwise vortices appearing in the near wake, which cannot be captured in a two-dimensional computation. The two-dimensional computations by Henderson [24] suffers the same drawback as well. Previous three dimensional simulations using the same discretization and solution techniques as described herein have shown good correlation with experimentally determined Strouhal numbers at higher Reynolds numbers [25].

The variation of C_L with time is shown in Fig. 4. A plot of the density contours at an intermediate time is shown in Fig. 5.

In order to determine the temporal order of accuracy, the test problem was solved using the same initial condition but with different time steps. The time interval of the study was approximately $1\frac{1}{4}$ shedding cycles. The solution at the end of the time interval is assumed to have accumulated the temporal error. Integral measures such as lift on the body, drag due to pressure forces and pitching moment of the body were then compared as follows to determine the order of accuracy. Let $G_{\Delta t}$ denote the integral measure being compared using a time step Δt , while G_{exact} denotes the exact solution. We do not know G_{exact} but based on the order of accuracy n of the scheme, we expect the following behavior:




Fig. 4. Variation of C_L with nondimensional time.




Fig. 5. Density contours calculated using a time step of 0.025.

$$G_{\Delta t} = G_{\text{exact}} + C_1(\Delta t)^n + \text{Higher order terms}, \tag{29}$$

where C_1 is a constant. By subtracting from Eq. (29) a similar expression for $G_{\Delta t/2}$ and neglecting the higher-order terms, we can obtain the following relation:

$$G_{\Delta t} - G_{\Delta t/2} = C_1 \left(1 + \frac{1}{2^n} \right) (\Delta t)^n = C_2 (\Delta t)^n, \tag{30}$$

where C_2 is another constant. Eq. (30) can be used to determine the order of accuracy of the scheme which can then be compared with the expected order of accuracy based on theory.

The order of accuracy was verified for two ESDIRK schemes (RK64 and RK43) and the second-order BDF (BDF2) scheme. The nonlinear systems which arose were converged until the maximum density correction, $|\Delta\rho|_{\text{max}} < 10^{-10}$. This ensures that the “iteration error” is negligibly small relative to the discretization error. Figs. 6–8 show the detailed refinement study for the RK64, RK43 and BDF2 schemes,




Fig. 6. Verifying order of accuracy of RK64 in $Re = 1200$ test case.




Fig. 7. Verifying order of accuracy of RK43 in $Re = 1200$ test case.




Fig. 8. Verifying order of accuracy of BDF2 in $Re = 1200$ test case.

respectively. It can be seen that all the integral measures yield nearly the same quantitative conclusions. The anomalous behavior for large time steps in Fig. 8 is likely due to the time steps being outside the asymptotic range. It is also seen that the computed order of accuracy is close to the expected order of accuracy. For example, Fig. 6 shows that the computed order of accuracy for the RK64 scheme is 3.8938 while the expected order of accuracy is 4. The small differences between expected and observed accuracy may be attributable to higher-order terms of this asymptotic analysis.

4.2. $Re = 100$ case

Another test case, $Re = 100$ was chosen as the wake is more laminar and two-dimensional at this Reynolds number. Using $\Delta t = 0.025$, the Strouhal number, St , was calculated to be 0.163. This value was again compared with a \sqrt{Re} -formula fit from [23].

- “Numerical” laminar flow, two-dimensional computations by Henderson [24]: the curve fit is given by

$$St = \left(0.2731 - \frac{1.1129}{\sqrt{Re}} + \frac{0.4821}{Re} \right),$$

which yields $St = 0.165$ at $Re = 100$.

- Experimental studies in the laminar wake of a circular cylinder: when $Re < 180$, the curve fit to experimental data is given by

$$St = \left(0.2665 - \frac{1.018}{\sqrt{Re}} \right).$$

This gives a $St = 0.1647$ at $Re = 100$.

It can be seen that our results compare well with both experimental results and previous numerical calculations.

In a manner, similar to the procedure described in the $Re = 1200$ test case, the order of accuracy of RK64 and BDF2 were verified as well. These results are shown in Figs. 9 and 10. It is again noted that the computed order of accuracy is very close to the expected order of accuracy for the different schemes.




Fig. 9. Verifying order of accuracy of RK64 in $Re = 100$ test case.




Fig. 10. Verifying order of accuracy of BDF2 in $Re = 100$ test case.

5. Parameter selection for solution schemes

The various solution schemes studied in this work contain a number of parameters which must be chosen optimally in order to fully exploit the potential efficiencies of each scheme. Because the LMG and PGMRES methods involve multiple levels of iteration nesting, (e.g. outer Newton, inner multigrid), these schemes contain more parameter choices and their overall performance is more dependant on the selection of these parameters. The NMG scheme has been used extensively in the past for steady and unsteady calculations and is well optimized for such calculations [26,5]. Furthermore, since all schemes rely on a common multigrid strategy, we choose to adopt a common multigrid approach based on the NMG experience, using four multigrid levels with a W-cycle for all schemes.

5.1. Parameter selection in the LMG method

The inexact Newton methods contain various parameters which must be chosen judiciously in order to optimize the overall run-time of these methods. For the LMG scheme, the parameters to be chosen include:

1. Number of LMG cycles carried out in $\mathcal{P}_{\mathcal{N}}$.
2. Number of smoothing iterations carried out on each grid of the multigrid.

Increasing either of these parameters would make $\mathcal{P}_{\mathcal{N}}$ a better approximation to $\tilde{\mathcal{A}}^{-1}$. Fig. 11 illustrates the effect of increasing the number of LMG cycles used in $\mathcal{P}_{\mathcal{N}}$. All the results shown in this section and the next are carried out using the BDF2 physical time-stepping scheme and a time step of $\Delta t = 0.05$.

We observe the following:

1. The rate of convergence slows down as the nonlinear residual decreases. This is due to the inexact solution of the linear system described by Eq. (20) at each Newton iteration, as we use only a single Richardson's iteration to solve the equation. Hence, the quadratic convergence expected of exact Newton's methods is not achieved.
2. Increasing the number of LMG cycles in $\mathcal{P}_{\mathcal{N}}$ can increase the nonlinear convergence by only a finite amount, and eventually asymptotes to a maximum rate. This is due to the fact the preconditioner $\mathcal{P}_{\mathcal{N}}$ is based on a first-order Jacobian, which even if inverted exactly does not correspond to the inverse of the exact Jacobian, A .

Since, the ultimate goal is to minimize the runtime required to converge the equations, $\mathbf{R}(\mathbf{w}) = \mathbf{0}$, to a given tolerance level, we plot the reduction of the residual against runtime for different choices of the number of LMG cycles. It can be seen from Fig. 12 that the use of two LMG cycles in $\mathcal{P}_{\mathcal{N}}$ is computationally most efficient for the range of error tolerances to which the nonlinear equations are converged.

Figs. 13 and 14 show the effect of performing different number of smoothing iterations on each grid of the multigrid. As expected, Fig. 13 shows that the number of Newton iterations required for a given level of residual reduction decreases with increase in the number of smoothing iterations.

The study indicates that the use of five smoothing cycles on each grid, while using two LMG cycles, results in a computationally optimal preconditioner, $\mathcal{P}_{\mathcal{N}}$, for the problem under consideration.




Fig. 11. Effect of changing the number of LMG cycles in $\mathcal{P}_{\mathcal{N}}$ of LMG.

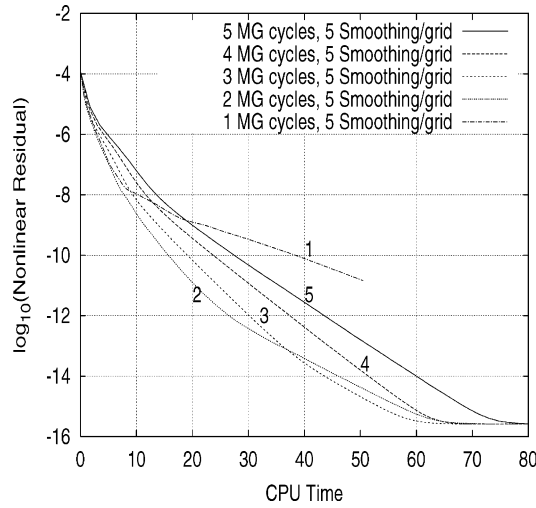


Fig. 12. Effect of changing the number of LMG cycles in $\mathcal{P}_{\mathcal{N}}$ of LMG on the runtime.

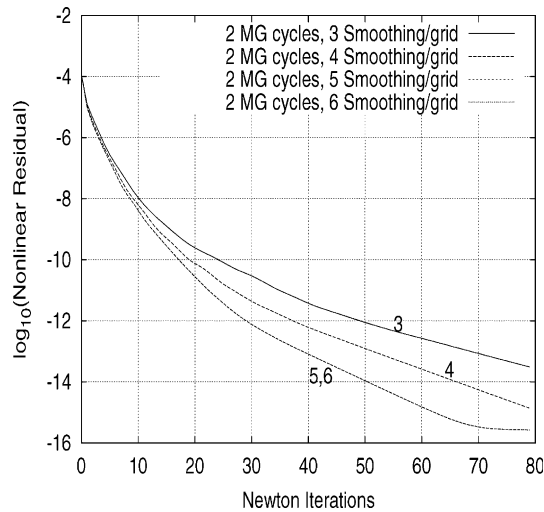


Fig. 13. Effect of changing the number of smoothing iterations carried out on each grid of the multigrid.

6. Parameter selection in PGMRES

Having selected all required parameters for the LMG preconditioner, $\mathcal{P}_{\mathcal{N}}$, the remaining parameter to be determined in the PGMRES method is the number of search directions. Increasing the number of search directions provides a more accurate solution to the linear system, which arises at every step of the Newton’s method. Thus, the Inexact Newton’s method approaches the exact Newton’s method as the number of search directions is increased.

Fig. 15 shows the effect of increasing the number of search directions on the number of Newton iterations required to converge to the solution of $\mathbf{R}(\mathbf{w}) = \mathbf{0}$. It can be seen that the effect of increasing the search directions is more pronounced as the nonlinear residual becomes smaller.

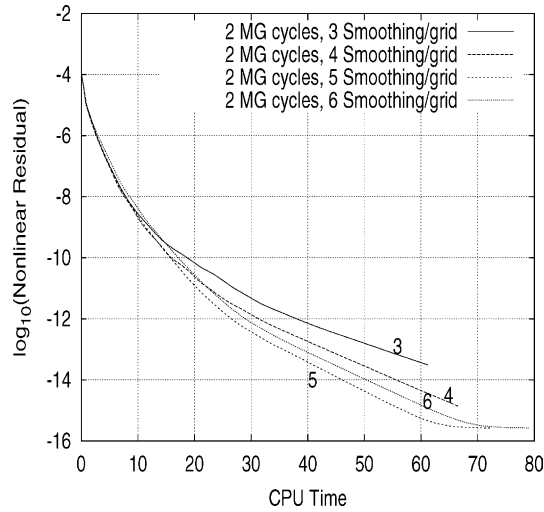


Fig. 14. Effect of changing the number of smoothing iterations carried out on each grid of the multigrid on the runtime.

However, increasing the number of search directions by one incurs the following additional computational overhead:

- A single evaluation of the nonlinear residual on the fine grid.
- A single evaluation of the preconditioner, $\mathcal{P}_{\mathcal{N}}$.
- Additional matrix–vector and vector–vector products required to compute the extra search direction and the optimal solution, $\mathbf{x}^{(m)}$ in a larger Krylov subspace.
- Additional storage for the extra search direction.

Hence, the choice of the number of search directions was decided based on minimizing the CPU time required for a given level of residual reduction. Fig. 16 plots the variation of the nonlinear residual against

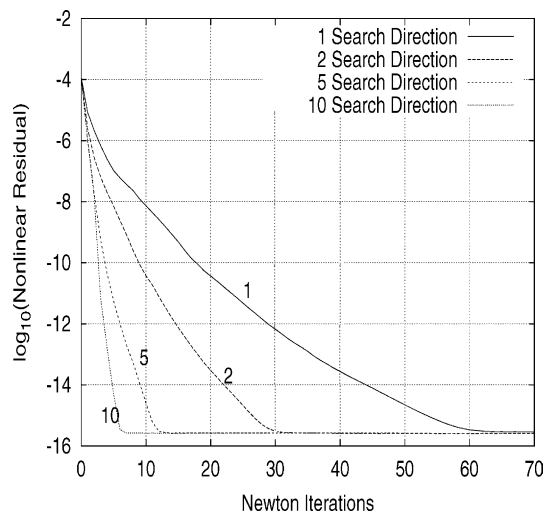


Fig. 15. Effect of changing the number of search directions in PGMRES.

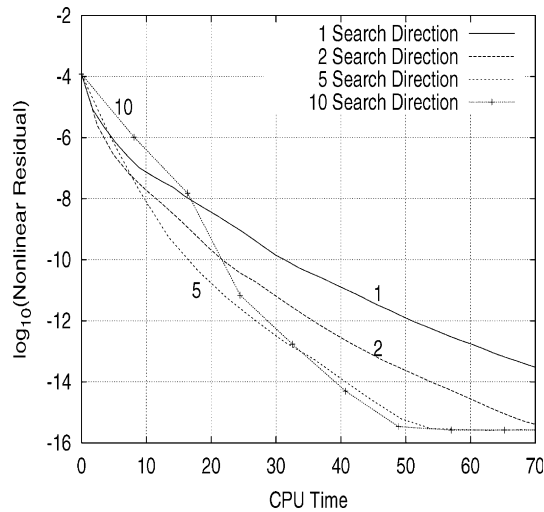


Fig. 16. Effect of changing the number of search directions in PGMRES on the runtime.

runtime for different choices of the number of search directions. It was determined that the use of five search directions was nearly optimal for most cases.

7. CPU time comparison of different schemes

In this section we examine the computational efficiency of two time-discretization schemes, RK64 and BDF2 as a function of accuracy levels. We simultaneously investigate the relative performances of the different implicit solution techniques discussed in this paper. Finally, we show that the combined improvements in efficiency obtained by using higher-order schemes and better nonlinear solvers such as LMG can result in up to an order of magnitude improvement in overall solution efficiency.

In order to compare the different schemes, we compare the CPU times required to advance the solution from an initial time T_i to a specified final physical time T_f , given an error tolerance in the final solution. In this study, we assume the error in the lift coefficient, C_L , to be a good measure of the integral error in the final solution. To measure the error in C_L , the numerical solution obtained using the RK64 scheme and $\Delta t = 0.0125$ was assumed to be numerically accurate with zero error. Finally, we choose three error tolerances, 10^{-2} , 10^{-3} and 10^{-4} , which we deem to be representative of engineering error tolerances, and make a detailed comparison of the different schemes. We choose a time interval $T_f - T_i = 1$, noting that the ratios of the CPU times of the various schemes should remain invariant with arbitrary choices of the time interval.

The physical time step, Δt , chosen to advance the equations to T_f depends on:

1. The physical time-stepping scheme, in this case either BDF2 or RK64.
2. The error tolerance level.

Fig. 17 plots the variation of the absolute error in C_L for the two different physical time-stepping schemes at $Re = 1200$. In order to obtain Fig. 17 the nonlinear equations at each time step were converged such that the rms error in the density residual was less than 10^{-10} , in order to avoid any contamination of the temporal error. The circle symbols in the figure indicate the time steps used in the following study for both schemes to achieve the three different prescribed error tolerances.

It can be seen that for the range of error tolerances considered, BDF2 requires a much smaller time step than RK64. Furthermore, as RK64 is fourth-order accurate in time, the error decreases faster with a decrease in time step, making it more efficient at lower error tolerance levels.




Fig. 17. Comparison of time steps required for BDF2 and RK64 schemes at $Re = 1200$.

Hence, the choice of the physical time-stepping scheme affects the overall efficiency of the simulation in three ways:

1. The number of time steps required to integrate to T_f , for a given temporal accuracy.
2. The total number of nonlinear solves per time step.
3. The condition number of the nonlinear systems to be solved. The nonlinear systems produced by the RK scheme are generally more difficult to converge due to the larger time step involved with the higher-order scheme.

Note that the explicit time step limit, (i.e. $CFL = 1$) is given by $\Delta t = 2 \times 10^{-4}$, since the minimum grid cell size in this case is 10^{-3} , and the Mach number is 0.2. Hence, for the high-accuracy cases using BDF2, the explicit time-step limit is approached. This accounts for the speedup in convergence of the implicit system for these cases. In fact, since multiple iterations are required to solve the pseudo-time problem, explicit time-stepping at the somewhat lower explicit time-step limit may be more efficient overall for these high accuracy cases. However, as the grid is refined (higher spatial resolution), or for larger time steps (lower temporal accuracy), the implicit schemes become most efficient.

We now try to quantify the efficiency gains by using higher-order schemes and a more efficient implicit solution technique. A major contributor to the inefficiency of implicit methods is solving the nonlinear systems at each stage/step to inappropriate sub-iteration tolerances. If the nonlinear system is solved to tighter tolerances, the additional work does not increase the overall solution accuracy. We do not attempt to answer this question in any detail, but assume that it is sufficient to converge the nonlinear residual to six orders of magnitude less than the error in C_L . This level was determined through numerical experiments, by setting various convergence levels and measuring the final temporal error. While the ratio of residual convergence to error in C_L appears relatively large, this is an artifact of the fact that these quantities are inherently scaled in different manners, since C_L represents a global integrated quantity, as opposed to the average flow field residuals.

In Table 2 we present the CPU times for the different combination of schemes.

In Figs. 18 and 19 we show the nonlinear convergence characteristics of BDF2 and the second stage of RK64 for the three different C_L error tolerances considered. Note that the required levels of convergence are more stringent for the smaller time steps, as shown by the bar symbol on each line. In the actual computations, the residuals were converged only to these levels. In these figures, however, convergence down to

Table 2

CPU times (in seconds on Pentium IV, 1700 MHz) required for carrying out one physical time step using the different nonlinear solvers, where time steps are chosen based on the given absolute error tolerances on C_L

Scheme	Δt	NMG	LMG	PGMRES
$C_L \text{ error} = 10^{-2}$				
BDF	0.01842	15.55	5.24	9.00
RK	0.24193	121.56	39.52	
$C_L \text{ error} = 10^{-3}$				
BDF	0.005825	12.45	4.59	9.01
RK	0.122047	131.4	43.84	
$C_L \text{ error} = 10^{-4}$				
BDF	0.001846	12.49	5.54	8.76
RK	0.067285	147.14	44.91	




Fig. 18. Comparison of the convergence of the nonlinear residual of BDF2 schemes for a single time step using NMG, where Δt 's are computed based on the given error tolerances on C_L . Bars denote levels of convergence required as a function of final C_L accuracy.

machine accuracy is shown to illustrate the overall convergence behavior. These examples all utilize the NMG scheme.

We observe the following:

1. Faster convergence for smaller error tolerances, as Δt is reduced.
2. Slower overall convergence for RK64 relative to BDF2 for similar error tolerances, as larger time steps are used in RK64.
3. Although not shown, similar behavior is observed for the other two methods (LMG, PGMRES) as well.

In Tables 3 and 4 we quantify the efficiency gains obtained by shifting from BDF2 to RK64 for the two different nonlinear solvers, NMG and LMG, respectively. The numbers in the different columns are ratios of the corresponding quantities used in RK and BDF respectively. The CPU time for a given time step is an order of magnitude larger for the RK schemes, and this ratio varies slowly with the increase in accuracy. This result is a function of the number of nonlinear solutions required per time step, the relative stiffness of




Fig. 19. Comparison of the convergence of the nonlinear residual of RK64 schemes for the first stage of a single time step using NMG, where Δt 's are computed based on the given error tolerances on C_L . Bars denote levels of convergence required as a function of final C_L accuracy.

Table 3
BDF2 to RK64 speedup factors for NMG

Abs. C_L error	$\frac{\Delta_{RK}}{\Delta_{BDF}}$	$\frac{CPU\ time_{BDF}}{CPU\ time_{RK}}$ for one time step	Overall gain
10^{-2}	13.13	0.128	1.68
10^{-3}	20.95	0.095	1.99
10^{-4}	36.45	0.085	3.1

Table 4
BDF2 to RK64 speedup factors for LMG

Abs. C_L error	$\frac{\Delta_{RK}}{\Delta_{BDF}}$	$\frac{CPU\ time_{BDF}}{CPU\ time_{RK}}$ for one time step	Overall gain
10^{-2}	13.13	0.133	1.75
10^{-3}	20.95	0.105	2.2
10^{-4}	36.45	0.123	4.5

each nonlinear problem, and the degree to which each system must be converged to maintain overall accuracy. However, the ratios of time steps between the two schemes is large and increases rapidly for higher accuracy levels, thus making the RK scheme more efficient overall, particularly at the more stringent error tolerances.

In Fig. 20 we plot the variation of the nonlinear residual for a single time step of BDF2 using different nonlinear solvers as a function of the number of Newton iterations for LMG and PGMRES, and as a function of NMG cycles in the case of NMG. The inexact Newton methods exhibit faster convergence rates per iteration than the NMG method. Furthermore, PGMRES is seen to outperform LMG as well, due to the fact that PGMRES produces a more accurate solution of the preconditioned linear system, Eq. (23) at each Newton iteration.

In Fig. 21 we plot the variation of the nonlinear residual for a single time step of BDF2 using different nonlinear solvers as a function of CPU time. We make the following observations:

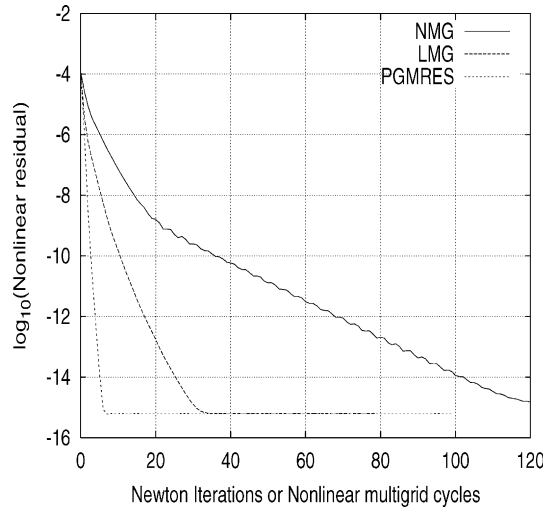


Fig. 20. Comparison of the convergence of the nonlinear residual of BDF2 schemes using the three different nonlinear solvers.

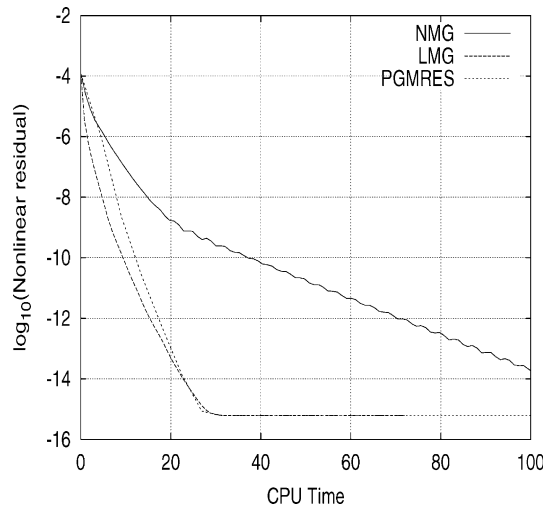


Fig. 21. Comparison of the convergence of the nonlinear residual of BDF2 schemes using the three different nonlinear solvers against CPU time.

1. LMG is computationally more efficient than NMG as it performs fewer nonlinear residual evaluations.
2. Although PGMRES has faster convergence in terms of Newton Iterations, the additional cost of each Newton iteration outweighs the gain in nonlinear convergence.

Finally, in Fig. 22 we show the CPU times required to reach $T_f - T_i = 1$ using the different schemes, versus the error tolerance level in C_1 . Fig. 22 clearly indicates that the commonly used combination of BDF2 and NMG exhibits the lowest computational efficiency for this problem. The most efficient solution strategy is obtained using the RK64 temporal discretization with the LMG solver, and these gains increase for higher accuracy levels.

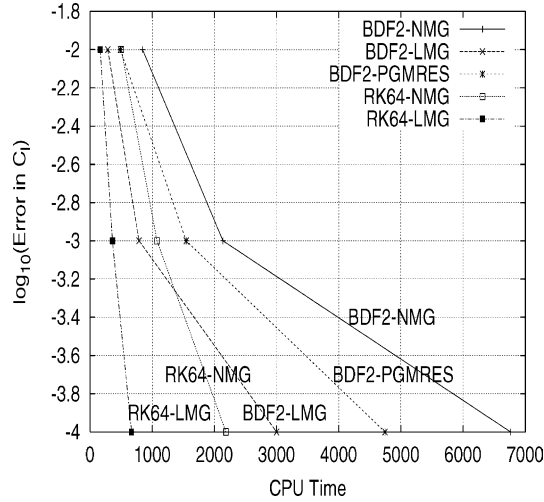


Fig. 22. CPU times required to reach $T_i = 1$ using the different schemes plotted against the error tolerance level in C_L .

Table 5
BDF2-NMG to RK64-LMG speedup factors

Abs. C_L error	$\frac{\Delta_{RK}}{\Delta_{BDF}}$	$\frac{CPU\ time_{BDF}}{CPU\ time_{RK}}$ for one time step	Overall gain
10^{-2}	13.13	0.393	5.16
10^{-3}	20.95	0.284	5.95
10^{-4}	36.45	0.278	10.14

In Table 5, we emphasize the overall gains obtained in shifting from the BDF2-NMG to the RK64-LMG solution strategy, as a function of the temporal accuracy. The gains obtained between the RK scheme and the BDF scheme and the gains obtained between the various nonlinear solution strategies are multiplicative, producing a larger overall gain than for either method used alone. This compounded efficiency gain increases with solution accuracy, yielding up to an order of magnitude improvement for the highest accuracy considered.

8. Conclusions

For problems where the temporal scales are of the same size as the spatial scales, explicit methods remain the most efficient solution strategy. However, for cases where temporal scales are much larger than the spatial scales, implicit methods are most efficient. Higher-order temporal discretizations such as RK64 can be expected to deliver superior efficiency for problems with smooth temporal variations, with increasing benefits for higher temporal accuracy requirements. The key to the efficient implementation of these implicit schemes is the use of an efficient nonlinear solver. The multigrid based solvers studied in this work deliver rapid convergence rates which are relatively insensitive to the separation of time and spatial scales.

For the test problem considered in this paper, the implicit multi-stage RK schemes have been shown to produce higher accuracy at reduced cost as compared to BDF2. Additionally, inexact Newton solution strategies have been shown to be well suited for solving the nonlinear systems which arise from temporal discretization at each time step.

The efficiency gains of both approaches are multiplicative, resulting in large potential savings when both methods are used in tandem. The combination of RK64 with LMG method worked very well for the error tolerances considered. The PGMRES algorithm studied in this work provided the fastest asymptotic convergence among all methods but was found to be noncompetitive due to the slower initial convergence of the method when only partial convergence of the linear systems is required. In cases where more accurate linear system solutions are required, PGMRES may be more competitive. Overall efficiency of the time integration schemes is greatly affected by the degree to which the nonlinear systems at each time step are converged.

The levels of convergence adopted in this work were determined a posteriori, and are therefore not predictive. A more exact quantification of the required convergence levels will be required in order to construct efficient time-dependent solution strategies. Future work will also include the use of temporal error estimation techniques coupled with dynamically adaptive time-step selection.

Acknowledgements

Much of the work was carried out while the first author was in residence at ICASE, NASA Langley Research Center. The authors would also like to thank the information given by Bijl and Carpenter regarding the various time integration schemes.

Appendix A. Butcher coefficients for RK64

The Butcher coefficients for the RK64 scheme used in the paper are given in the table below. RK64 is a six stage with scheme, with fourth-order accuracy and five implicit stages. As in all ESDIRK schemes, $b_j = a_{6j}$:

0	0	0	0	0	0
0.25	0.25	0	0	0	0
0.137776	-0.055776	0.25	0	0	0
0.144636866	-0.2239319076	0.4492950416	0.25	0	0
0.098258783284	-0.59154424282	0.8102105383	0.28316440571	0.25	0
0.15791629516	0	0.18675894052	0.68056529531	-0.275240531	0.25

References

- [1] H. Bijl, M.H. Carpenter, V.N. Vatsa, Time integration schemes for the unsteady Navier–Stokes equations, AIAA Paper 2001-2612.
- [2] A. Jameson, Time-dependent calculations using multigrid, with applications to unsteady flows past airfoils and wings, AIAA Paper 91-1596.
- [3] N.D. Melson, M.D. Sanetrik, H.L. Atkins, Time-accurate Navier–Stokes calculations with multigrid acceleration, in: Proc. 6th Copper Mountain Conference on Multigrid Methods, NASA Conference Publication 3224, 1993, pp. 423–439.
- [4] V. Venkatakrisnan, Implicit schemes and parallel computing in unstructured grid CFD, in: VKI Lecture Series VKI-LS 1995-02, 1995.
- [5] D.J. Mavriplis, An assessment of linear versus non-linear multigrid methods for unstructured mesh solvers, Journal of Computational Physics 175 (1) (2002) 302–325.
- [6] D.J. Mavriplis, Directional agglomeration multigrid techniques for high-Reynolds number viscous flows, AIAA Journal 37 (10) (1999) 1222–1230.

- [7] G. Dahlquist, A special stability problem for linear multistep methods, BIT 3.
- [8] C.A. Kennedy, M.H. Carpenter, Additive Runge–Kutta schemes for convection–diffusion–reaction equations, Tech. Rep. NASA/TM-2001-211038, NASA Langley Research Center, 2001.
- [9] E. Hairer, S.P. Norsett, G. Wanner, Solving Ordinary Differential Equations I: Nonstiff Problems, Springer, Berlin, 1991.
- [10] V.A. Mousseau, D.A. Knoll, W.J. Rider, Physics-based preconditioning and the Newton–Krylov method for the non-equilibrium radiation diffusion, Journal of Computational Physics 160 (2000) 743–765.
- [11] E.J. Nielsen, W.K. Anderson, R.W. Walters, D.E. Keyes, Application of Newton–Krylov methodology to a three-dimensional Euler code, in: Proceedings of the 12th AIAA CFD Conference, no. AIAA Paper 97-1953, 1997.
- [12] Y. Saad, M.H. Schultz, A generalized minimal residual algorithm for solving non-symmetric linear systems, SIAM Journal on Scientific and Statistical Computing 7 (1986) 856.
- [13] D.J. Mavriplis, Multigrid strategies for viscous flow solvers on anisotropic unstructured meshes, Journal of Computational Physics 145 (1) (1998) 141–165.
- [14] D.J. Mavriplis, Multigrid techniques for unstructured meshes, in: VKI Lecture Series, no. 1995-02 in VKI-LS, 1995.
- [15] B.R. Hutchinson, G.D. Raithby, A multigrid method based on the additive correction strategy, Numerical Heat Transfer 9 (1986) 511–537.
- [16] B.V. Leer, C.H. Tai, K.G. Powell, Design of optimally-smoothing multistage schemes for the Euler equations, AIAA Paper 89-1933.
- [17] E. Turkel, Preconditioning-squared methods for multidimensional aerodynamics, in: Proceedings of the 13th AIAA CFD Conference, Snowmass, CO, June, 1997, pp. 856–866.
- [18] R. Dembo, S. Eisenstat, T. Steihaug, Inexact Newton methods, SIAM Journal on Numerical Analysis 19 (400).
- [19] J.W. Ruge, K. Stüben, Algebraic multigrid, in: S.F. McCormick (Ed.), Multigrid Methods, SIAM Frontiers in Applied Mathematics, SIAM, Philadelphia, 1987, pp. 73–131.
- [20] G.H. Golub, C.F.V. Loan, Matrix Computations, third ed., Johns Hopkins University Press, Baltimore, MD, 1996, pp. 508–513 (Chapter 10).
- [21] L.N. Trefethen, D. Bau, Numerical Linear Algebra, SIAM, Philadelphia, 1997 (Chapter 35).
- [22] L.B. Wigton, N.J. Yu, D.P. Young, GMRES acceleration of computational fluid dynamics codes, in: Proceedings of the 7th AIAA CFD Conference, 1985, pp. 67–74.
- [23] C. Williamson, G. Brown, A series in $1/\sqrt{Re}$ to represent the Strouhal–Reynolds number relationship of the cylinder wake, Journal of Fluids and Structures 12 (1998) 1073–1085.
- [24] R. Henderson, Nonlinear dynamics and pattern formation in turbulent wake transition, Journal of Fluid Mechanics 352 (1997) 65–112.
- [25] D.J. Mavriplis, J. Pelaez, O. Kandil, Large eddy and detached eddy simulations using an unstructured multigrid solver, in: C.L. Liu, L. Sakell, T. Beutner (Eds.), DNS/LES – Progress and Challenges. Proc. of the Third AFOSR Int. Conf. on DNS/LES, 2001, pp. 461–470.
- [26] D.J. Mavriplis, On convergence acceleration techniques for unstructured meshes, AIAA Paper 98-2966, Presented at the 29th AIAA Fluid Dynamics Conference, Albuquerque, NM, June 1998.